2 CMP Technology

Michael R. Oliver

2.1 Background and Motivation for CMP

Chemical Mechanical Polishing, also often referred to as Chemical Mechanical Planarization (CMP), was initially used as an enabling technology to fabricate high performance multiple level metal structures. Specifically, after the first level of metal was fabricated, and a nearly conformal silicon dioxide interlevel dielectric (ILD) layer was deposited, the second level metal has several fabrication problems, including deposition, resist patterning and etching. These difficulties are caused by the steps in the topography over which this layer must be processed [1]. Other technologies, especially spin-on glass (SOG), reduce many of the problems of multi-level metal integration approaches, however SOG introduces additional difficulties of its own, and has been primarily used for two and three level metal structures [2].

From a technology point of view, the initial work to develop CMP for semiconductor fabrication was done at IBM [3], where they used the expertise of their own silicon wafer fabrication technology. This expertise included an understanding of the hardware: machines, pads, and slurries. The scientific understanding of CMP was largely based on that of glass polishing [4], but that theory itself was not quantitative. In 1990, Cook presented an excellent summary of the understanding of the mechanisms of glass polishing up to that date [5]. He emphasized the poor quantitative agreement of existing models with experimental results.

Once the technology and the required equipment were available, the application of CMP quickly spread beyond polishing inter-level dielectric (ILD) layers. For example, CMP began to be used instead of reactive ion etching (RIE) to remove tungsten which was deposited to fill the via openings between metal layers [6]. Another metal CMP application, the fabrication of inlaid trenches filled with metal, also called damascene, was proposed. This polishing technology has essentially been an enabling technology for the introduction of copper interconnects into standard semiconductor processing. Until the availability of CMP, copper was not used for interconnects even though it has a lower resistivity than aluminum for the reason that it could not be easily be etched by RIE [7].

Other applications for CMP have also emerged. A very significant one is

Shallow trench isolation is an integration approach that allows transistors to be packed at a higher density by reducing the isolation spacing between adjacent transistors. Another use for CMP is polishing polysilicon via plugs and capacitor structures in memory devices [10].

The purpose of this volume is to describe the major applications of CMP in the current semiconductor technology. Broadly speaking, CMP technology can be divided into two areas, dielectric and polysilicon CMP and metal CMP. Oxide CMP, which is the polishing of silicon dioxide, will be used for this chapter as the vehicle to discuss the elements of CMP. This chapter will also cover the technology and application of other dielectric polishing applications.

The first section describes the elements of the CMP process. Development and refinements of the basic approach will follow. The last part of the chapter will address dielectrics other than silicon dioxide.

2.2 Description of the CMP Process

In the current standard approach, Chemical Mechanical Polishing takes place where the surface of the wafer to be polished is forced against a polishing pad. The polishing pad is covered with a liquid slurry which contains abrasive particles. The wafer is moved relative to the slurry-covered pad, and the rate at which material is removed from the wafer is often described by the heuristic equation called Preston's Law [11]:

$$RR = K_p * P * V \tag{2.1}$$

with

RR – removal rate

K_p – a constant, Preston's coefficient

P – local pressure on wafer surface

V – relative velocity of the point on the surface of wafer vs. the pad.

This relationship is empirical, a system where material was removed by grinding. Numerous dielectric and metal CMP models have been, and are continuing to be, proposed in the literature, and for most, Preston's Law is only an approximation. However, for much of the data obtained in practice, especially silicon dioxide CMP, Prestons Law provides a reasonably good fit to the data.

2.3 Polishing Equipment

The first polishing machines on which semiconductor CMP processes were developed were rotary polishing tables. As the machine technology has ad-

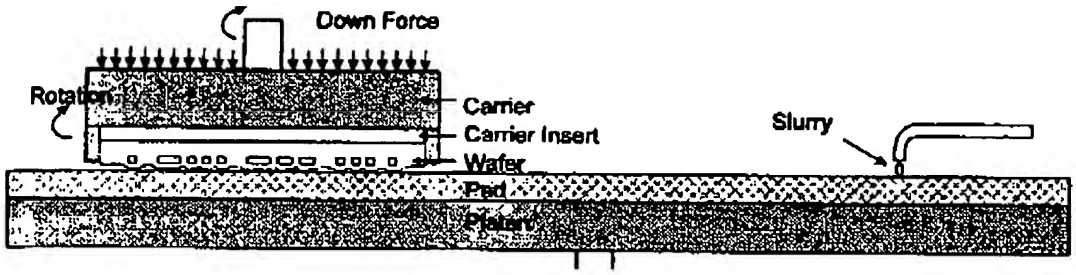

Fig. 2.1. Drawing of basic rotary CMP machine, showing wafer, carrier and platen (table). From US Patent 4,944,836. The retaining ring holds the wafer under the carrier insert (pad) (see text)

vanced, machine designs have evolved and other basic designs have been employed as well. However, most of the machines currently being sold as well as those in use are rotary tools.

A representative rotary polishing machine is diagrammed in Fig. 2.1, which is from [3], one of the early IBM patents. In such a machine, the polishing pad is circular and the wafer is placed in a carrier face down and is forced against the pad while the pad table, or platen, is rotated on its own axis.

The forces applied through the carrier on to the wafer are generally in the range of 1–10 psi, with oxide polishing usually in the higher end of the range and metal polishing in the lower end of the range. In practice, the table diameter is in the 20–26" range for commercial CMP machines which typically polish one wafer at a time. Figure 2.1 shows just the simplest configuration for a single table, single head (wafer carrier) system. In Chap. 5, Thomas Tucker reviews with details the many options for rotary designs as well as other designs. A key element of any polishing machine is to have well controlled pressures applied uniformly over the wafer as well as having controlled table and carrier rotation rates.

There are several other features in Fig. 2.1 that are to be noted. One is that slurry is dispensed from a tube in front of the wafer, so that as the table rotates, it is pulled under the wafer. Also, though not easily visible on this scale, the retaining ring around the edge of the wafer keeps the wafer in the carrier. The bottom of the retaining ring is recessed, usually about 0.008", from the plane of the bottom of the wafer.

The conditioner is a mechanism that moves a hard abrading surface, often a matrix with embedded diamond points, across the pad surface to roughen it. This is critical to CMP as an inadequately roughened pad surface results in a very low polish rate [13].

The slurry that flows onto the pad covers the roughened pad surface which moves under the wafer. The grooves on the pad allow more slurry to be brought under the retaining ring to the wafer face. As is discussed in Chap. 6, many pad structures also have small hollow spherical pores that are

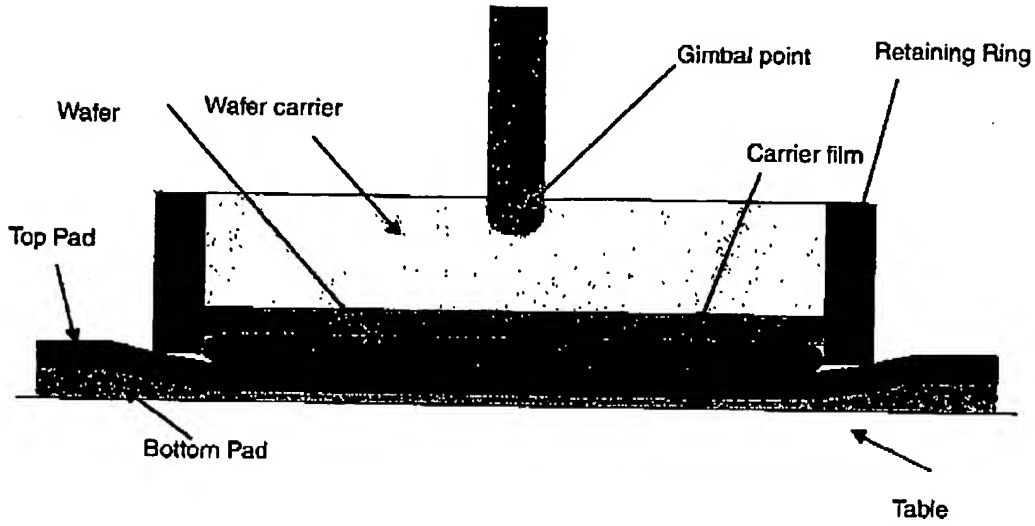

Fig. 2.2. Typical wafer carrier cross section, not to scale (see text)

exposed to the surface. These also contain slurry that is brought to the face of the wafer as the table rotates.

A second view, showing a simplified cross-section of a representative carrier, is shown in Fig. 2.2. This also shows a two-layer polishing pad, as well as the carrier film behind the wafer. Key features include the application of the down force from the carrier arm to the carrier at the gimbal point. The body of the carrier rotates about the gimbal point. Note that the gimbal point is above the wafer. The bottom pad layer and the carrier film are relatively compressible. Both films are generally about 0.050" thick and each compresses about 2–4% at pressures in the 5–7 psi range. The reason that both of these relatively compressible films are used is to maintain, over the entire wafer, a nearly uniform pressure at the wafer-polishing pad interface within the variations of the pad thicknesses, wafer thickness and the dimensional control of the table flatness relative to the carrier. It is worth noting that as machine and process tolerances become tighter, the sub-pad and carrier films can be thinner, since they will not have to compensate for as much mechanical variation.

Because the gimbal position (for most gimbal carrier designs) is about 1" above the wafer-pad interface, when the table rotates the friction at the wafer-pad interface causes a moment about the gimbal point, which increases the downward pressure at the leading edge of the wafer. Since the total constant downward force is applied at the gimbal point, a locally higher pressure at the leading edge of the wafer will also create a reduced pressure at the trailing edge. The exact instantaneous local pressures across the wafer will depend on properties of many of the elements in the system. One of the purposes of carrier rotation is to average out the leading and trailing edge effects [14]. This rotation averages the locally high removal rates at the leading edge

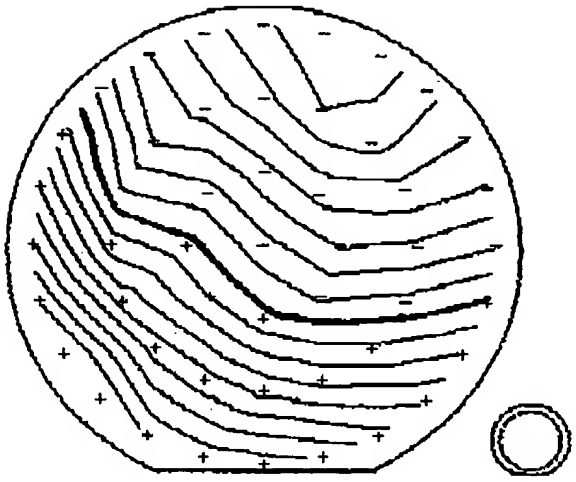

Fig. 2.3. Wafer polished for 60 seconds on Strasbaugh 6DS with no carrier rotation. Pre-polish wafer thickness was 10,000 Å and the *dark line* is the post-polish 7500 Å contour line. The *contour line* spacing is 250 Å. The leading edge is at the top of the wafer. Courtesy of David Evans, private communication

and the correspondingly low rates at the trailing edge, and can substantially reduce the non-uniform polish rate observed with a stationary carrier.

An example of such a polish rate variation is shown in Fig. 2.3. There for typical conditions except for no carrier rotation, the leading edge of the wafer has a higher polish rate than the trailing edge. There also is an effect of the outer side (here the right side) of the wafer polishes more quickly than the inside as it has a higher linear velocity and there is no velocity averaging by table rotation. At the trailing edge, there is less than 500 Å/min polish rate, and at the leading edge the rate is greater than 4250 Å/min.

Polishing conditions:	ILD1300 silica based slurry
	IC1400 perforated pad
	Down force – 9 psi
	Table rotation rate – 40 rpm.

However, with carrier rotation and in the absence of a wafer flat, or any other significant departure from rotational symmetry, the polish rate, and the total amount removed, will have close to radial symmetry. This behavior is widely observed on machines with rotating carriers.

A compressible carrier film between the carrier and the wafer is required to help provide a nearly uniform force on the back of the wafer with the variations in wafer thickness and top polish pad thickness. This is especially important as the pad wears with use. As discussed in Chap. 5, a trough is formed in the pad in the wafer path through abrasion during polishing. This trough can be quite deep, up to 25 μm or more lower than the edges of

and minimized in most current polishing systems by varying the conditioning conditions, primarily dwell time, as a function of radial position on the pad [15]. In general, though, some radial variation of pad thickness usually exists. Also, there are thickness non-uniformities due to manufacture of the polish pad as well as a lack of true planarity of the polishing table.

The compressible bottom pad, which is usually an impregnated felt or foam, is another component introduced to maintain a nearly constant pressure on the bottom side of the hard urethane polish pad (see Chap. 6). Because the top pad is less stiff than the wafer, the two layer stack of the urethane polish pad and the softer bottom pad determine key polishing features when polishing wafers with topography, i.e., wafers with device structures. (see Planarization section below and Chap. 10).

In summary, the purpose of the machine and pad elements together, is to provide as uniform polish conditions (pressure, velocity) as possible at all points on the wafer. The system is also designed to provide removal rate averaging through table and carrier rotation to minimize total variation of the amount of silicon dioxide, or other polished film, remaining across the wafer.

2.4 Polish Process

Several key elements of CMP are worth emphasizing. CMP of silicon dioxide surfaces requires certain specific properties of the slurries and pads. The slurries require the use of certain metal oxides as the abrasive particles. The oxide most widely used is silica (silicon dioxide), which can be fabricated by various methods (see Chap. 7). However, other metal oxide particles, such as ceria and manganese dioxide, can also be used. The slurry liquid needs to be aqueous. For maximum polish rates with silica slurry, the pH of the slurry should be in or near the range of 10.5–11.2. In this regime, the surface of the silicon dioxide film is strongly hydroxylated with internal bonds broken by interaction with the alkaline liquid. However, if the pH is much greater than 11.5, the silicon dioxide film will break down entirely and simply begin to dissolve [17].

There is a wide range of silica particle size that is used for oxide CMP. Mean diameters range from about 25 nm for some colloidal silica slurries to about 300 nm for some fumed silica slurries.

The specific properties of the particles and the solutions are covered in detail in Chaps. 3 and 7. The use of other abrasive particles or other liquids generally results in little or no material removal, only some level of surface scratching.

There are several types of polishing pads used for silicon dioxide polishing. Softer pads, such as poromeric pads, are often used for local smoothing or scratch removal, also call buffing. But such pads have poor planarization

polished surfaces with longer range planarization (see Chap. 7). Though hard pads other than urethane pads are available, almost all silicon dioxide CMP is done with urethane pads. Urethane pads usually contain spherical pores or voids with diameters in the 30–50 μm range. These pores comprise about one third of the total pad volume, and also the same proportion of the top surface area. The surface of the pads can also be manufactured to have grooves or perforations. This surface pad structure aids slurry transport across the wafer surface.

The most significant feature of the urethane pads that is key to the CMP process is the formation of asperities on top of the pad by the process of conditioning. As noted, a typical conditioner has diamond points embedded in a matrix. This matrix is pressed against the pad while it is moving, and the conditioner is rotated. The surface of the pad is roughened to a level depending on the equipment and operating point. For representative conditions, in the space between the pore openings the pad surface has a roughness, R_a, of 1–5 μm with a spatial frequency of the same dimensions. An example of such a newly conditioned surface is shown in Fig. 2.4.

Empirically, the correct abrasive and liquid for the slurry, as well as pads with specific properties and appropriately conditioned surfaces are all required for the CMP process to occur. These lead to the working model pictured in Fig. 2.5 of how a specific film removal event takes place during CMP.

In Fig. 2.5, the silicon dioxide film is polished when an abrasive particle is forced against the film by an asperity of the pad. The particle is, under the force of the asperity pushing against the film, dragged along by the asperity at the relative velocity of the pad with respect to the wafer. However, the interaction of the particle with the film is not clearly understood. There recently have been proposed several alternative models of asperity-abrasive film

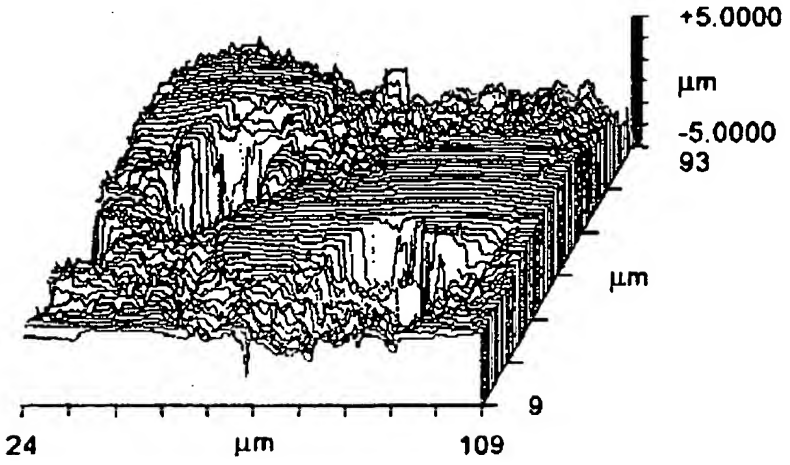

Fig. 2.4. Image of surface of conditioned IC1000 pad showing pores and conditioned surface over a $108 \times 144(\mu m)^2$ area. Image taken with Zygo NewView 500. Courtesy of Robert Schmidt, Rodel

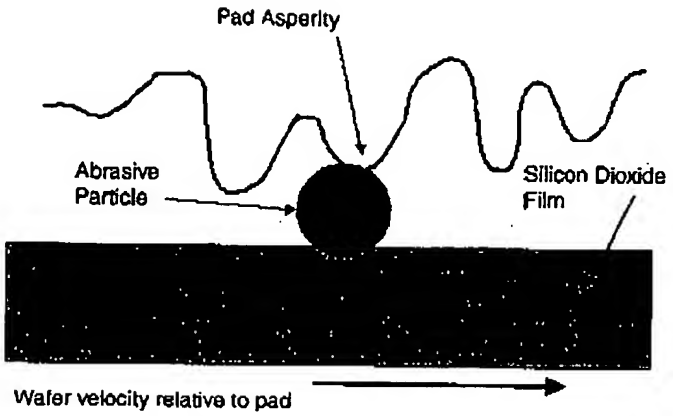

Fig. 2.5. Components (idealized) of film removal by CMP include the abrasive particle forced against the film by a pad asperity. The film to be polished moves relative the asperity with the abrasive

interaction that have led to overall relations between polish rate as various functions of pressure and velocity that do not follow Preston's Law [19, 20]. David Stein has compared several of these models to observed polishing data, especially in low pressure regimes [12], but none is an improvement over the simple Preston model. In an earlier work [5], Cook summarized the models for glass polishing to that date (1990), and much of this discussion is directly applicable to oxide CMP. Unfortunately, up to now there has been no clear quantitative, or semi-quantitive, interaction mechanism model proposed that is in reasonable agreement with the observed data.

The action of individual particles in polishing is repeated continually as the polishing process proceeds. As a result of carrier rotation, there is no preferred directionality for the paths of the polishing particles so that the sum of all the polishing events per unit time is the production of a average removal rate of the film.

2.5 Planarization

In contrast to glass polishing or silicon wafer polishing, for ILD polishing the goal of the CMP process is to planarize topography created by previous semiconductor processing steps. For other polishing steps such as STI CMP processes (see below) or metal CMP processes (see Chap. 3), CMP is used to remove an overburden of one material and stop on another material, leaving a planar surface.

In general, topographical features have different local polishing rates than do planar surfaces. Consider a polishing system where Preston's Law is a good approximation for the local polishing rate over a wide range of pressures, i.e.,

For the case where the velocity, V_0, and the term, K_P, are held constant, the local pressure determines the polish rate. This Prestonian relationship is generally valid for silica slurries polishing silicon dioxide. For the situation where the entire wafer has the same pattern density, then the local pressure on the top of each feature can be determined, as pictured in Fig. 2.6, as the average pressure applied to the pad divided by the pattern density, since the force per unit area applied to the pad is applied over the area of the pattern. For the case where the pattern density is some fraction of the total area, ϱ_1, then the down force is applied to this reduced area, and the polishing rate of each of the features will be increased to

$$RR_1 = K_P P_1 V_0 \qquad \text{here} \quad P_1 = P_0/\varrho_1. \tag{2.3}$$

Or here for the reduced density, ϱ_1,

$$RR_1 = K_P(P_0/\varrho_1)V_0. \tag{2.4}$$

For this sparse region where the feature density is uniform and at a density ϱ_1, the features will polish at a rate determined by the local Preston's Law. For example, if $F_1 = 0.25$, or 25%, then the features will polish at (1/0.25), or 4, times the rate of the planar surface.

The variation in local polish rate with feature density does not require a simple Preston's Law relationship between pressure and polish rate. If the polish rate on a planar surface can be described as

$$RR = f_A(P), \tag{2.5}$$

then for uniform features of density ϱ_K, the removal rate is described by

$$RR_K = f_A(P/\varrho_K). \tag{2.6}$$

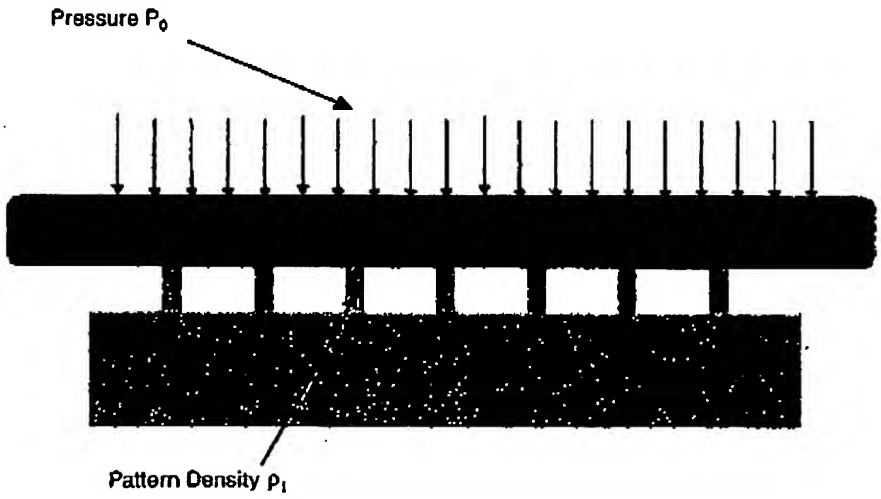

Fig. 2.6. Pressure P_0 is applied to pad and transmitted on to a wafer with a feature

Such a pressure dependence can and does occur in metal CMP and also in non-silica abrasive silicon dioxide polishing [21].

The effect of local polish rate dependence for patterned features has been studied by many workers. In Fig. 2.7 is pictured some data showing increased polish rate with decreasing density [22]. When the features are eliminated, the polish rate then reduces to the rate for a planar surface. The time to reach this transition to planarity decreases with decreasing density since the local polish rate increases with decreasing pattern density.

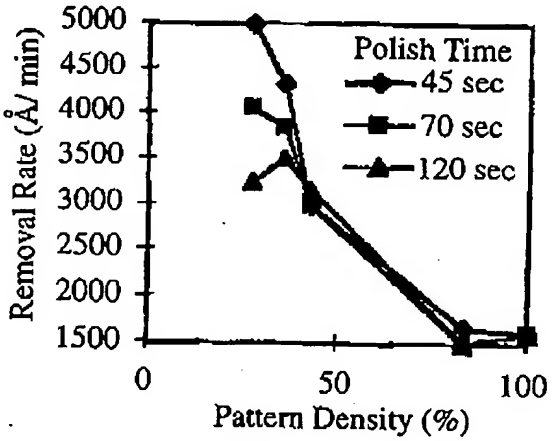

Fig. 2.7. Polish rate with a silica slurry as a function of silicon dioxide structure density. From [20]

In most CMP systems, the polishing pad is somewhat flexible. Also, in practice, there are variable pattern densities within a die and across the wafer. Thus the pattern density in the vicinity of a given point affects the local polish rate. It has been shown ([21] and Chap. 9) that a weighting function of the local pattern density out to a certain distance can effectively determine the local polish rate. The weighting function is a decreasing function with distance, so that features close to the local area of interest have the greatest effect on the polish rate. The range over which pattern features can affect one another is a function of the polishing system, primarily the thickness and modulus of the polishing pad. Though the range can be modeled to be a fixed length, with no influence beyond that length, the actual interaction decreases gradually. These and related issues are covered in depth in Chap. 9.

Areas with different local densities that are sufficiently separated will have independent polishing behaviors. Those areas with the lowest pattern densities will polish the most quickly, and those with the highest the most slowly. Once a given independent area is planarized, it will polish at the planar rate. This is pictured in Fig. 2.8 from [21], where low, medium and high density areas are pictured during stages of simultaneous polishing.

As seen in Fig. 2.8, once the entire wafer has been planarized, different

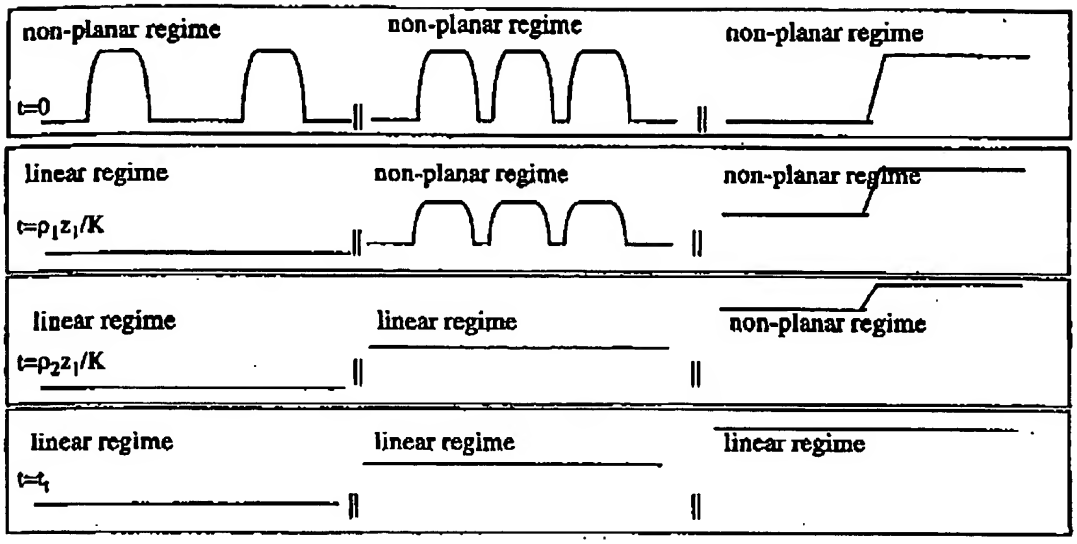

Fig. 2.8. Different final planarization thicknesses remain depending upon initial pattern density. See text. From [21]

film remaining. For adjacent areas the transition between these two areas will occur over a distance generally referred to as the planarization length. This is pictured in Fig. 2.9. The planarization length is a function of the interaction distance of the polishing system, and the amount of the polished film that has been removed during the polish step. Once local planarization is achieved as shown in Fig. 2.8, the planarization length will slowly grow as polishing continues. These lengths are generally hundreds of microns, and this subject also is discussed further in Chap. 9.

There are several consequences of the different clearing times and the resulting longer range thickness variations that exist once the topography has been removed. The first is that different wafers with different patterns will, in general, require different polish times to remove all of the topography. A very sparse metal pattern covered with a deposited ILD silicon dioxide layer

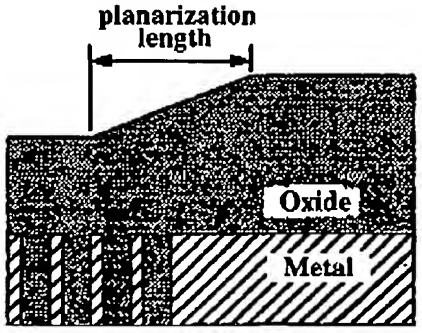

Fig. 2.9. Planarization length, the transition length between post-CMP high and low regions

will have its topography removed more quickly than will a wafer with a dense pattern. This well-known phenomenon often requires different polish recipes for the same process step for wafers with different patterns.

Another problem associated with variable feature clearing time is that, if one area of a chip has a sparse features such as the first pattern in Fig. 2.8 and another part has a dense pattern as does the third pattern, within die non-uniformity (WIDNU) will be large. This problem is a serious one. A frequently used approach to address this is to insert dummy structures so that all areas of the die will have similar feature densities. Dummy structures are isolated features that are designed to affect the CMP step and not have an electrical interaction [25]. In multi-level structures, the post-CMP surface of one level is the substrate upon which the metal and dielectric film of the next level are deposited. Since true planarity is not achieved at each level, the magnitude of the non-planarity can grow with multiple levels and is a key consideration in the integration process (see Chap. 10).

In addition to these long range planarization effects, there is a shorter range phenomenon that occurs when the height of the topography is reduced to the range of 400 nm or less. Ideally, no polishing at the bottom of a step in topography should occur until the step is removed. In practice, however, the bottom of the step begins to be polished before the step is removed [22] and the step height is not reduced at the ideal rate. As a result, in order to remove the step an extra amount of the film below the bottom of the original step must be removed. This additional amount of silicon dioxide that is deposited and then removed is another factor that must be accounted for in the overall integration considerations (see Chap. 10). Representative curves from [22] for various pattern densities are shown in Fig. 2.10. The ideal curves for each of the densities are compared to the data. It is, of course, desirable that the

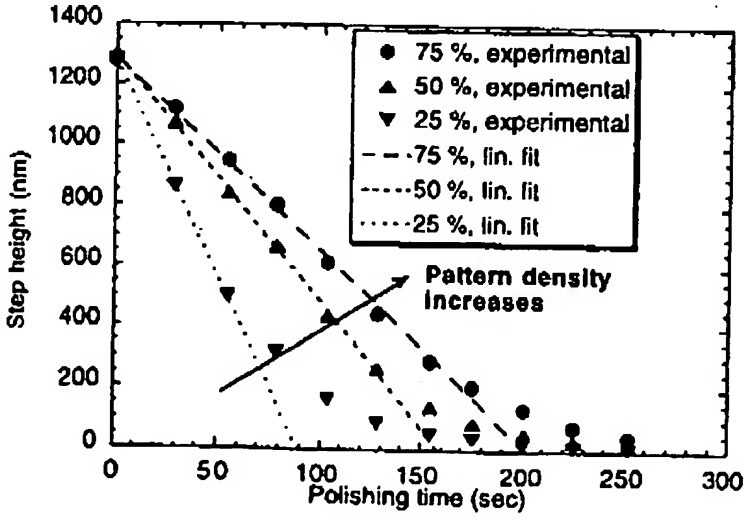

Fig. 2.10. Step height decrease vs. polish time for different density structures. Note departure from linear decrease below about 300 nm. From [22]

departure from the ideal be as small as possible. Slurryless, or fixed abrasive, technology, where the abrasive particles are embedded in the polishing pad, shows great potential for approaching such ideal planarization characteristics.

2.6 Polish Process Variables

The polishing removal rate, at least for silicon dioxide polishing on planar surfaces, is often well described by Preston's Law (2.1). This says that, for a given system, the removal rate is linearly proportional to local pressure and velocity between the pad and polished film. All of the other variables of the system are incorporated into the constant, K_p. These variables include the properties of slurry and pad, as well as the temperature. In addition, the properties of the material being polished are significant.

There is a very large range of possible system operating points, but in semiconductor fabrication, many of the components of the polishing systems are nearly standard across the industry, so variation of polishing performance within a narrow specific range is of most interest. However, as in other areas of technology, substantial changes in operating conditions are continually evaluated, and upon occasion, offer a significant advantage for some performance parameters, and then the new operating conditions (or equipment) are then adopted by a group of users.

Before discussing the variations of system parameters embedded in K_p, the ranges of pressure and velocity will be covered.

2.6.1 Pressure and Velocity Variation

For a given pad, slurry and polish film at ambient temperature, K_p, can be considered constant, and we can consider the pressure and velocity variations. For representative conditions, a typical slurry for polishing silicon dioxide contains 13 wt% solids of silica in a basic solution. For a standardly conditioned urethane polishing pad, a typical polish rate behavior as a function of average down force for a fixed table rotation frequency is shown in Fig. 2.11a. A corresponding curve for polish rate as a function of table rotation frequency for fixed average down force is shown in Fig. 11b. The film being polished is silicon dioxide deposited by plasma enhanced chemical vapor deposition (PECVD), which is a standard semiconductor deposition process. It can be seen that Preston's Law is in reasonable agreement with the data over the range tested, with some departure at very high table speeds.

In current practice, with pads and slurries like the above, average down forces on the wafer rarely exceed 10 psi. This is because the high total forces applied to the pad-wafer-slurry system result in the wafer not traveling smoothly over the pad surface, but sticking at points. This usually leads to wafer breakage or other forms of damage. As a result, with the current

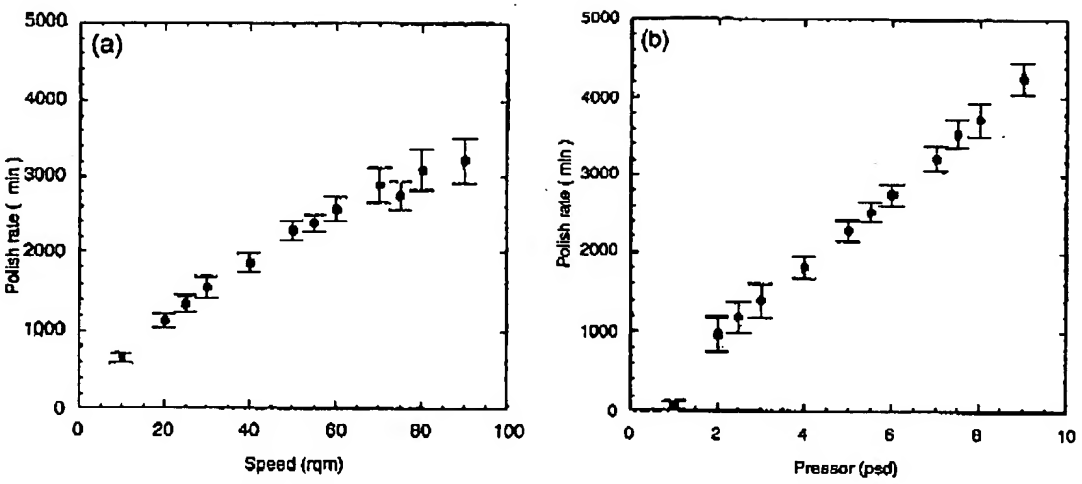

Fig. 2.11. (a) Polish rate vs. table rotation rate at 5 psi downforce on Westech 472 polisher. (b) Polish rate vs. downforce for 50 rpm table rotation rate on Westech 472 polisher. Both figures courtesy of David Stein

polishing machines, pads and slurries, semiconductor CMP processing is generally done below 10 psi. This, of course, may change over time with machine and pad design evolution. The table rotation rates shown in the two graphs of Fig. 2.11a and 2.11b are for a Westech 372M machine with an average radius position 16 centimeters from the center of the table. The magnitude of the instantaneous linear velocity of any point on the wafer is then given by

$$V = 2\pi r f, \tag{2.7}$$

where

V = magnitude of linear velocity, or speed, at that radius,

r = radius of the given point on the wafer with respect to the center of the table, and

f = rotation frequency of the table.

For this system at 60 rpm, or 1 rps, for the center of the wafer,

$$V = 1.01\,\text{m/s}. \tag{2.8}$$

There is a trend with machine improvements to design machines to operate at higher linear velocities in order to produce higher polish rates. Representative polish speeds for newer equipment designs are up to twice this speed or more. This issue is addressed in Chap. 5. The carrier rotation rate also affects the average speed at a given point on the wafer, and this effect increases as a function of the position on the wafer relative to the center of the wafer. Rotation of the carrier serves to average the polish direction over the entire wafer, but at very high carrier rotation rates, it may change significantly the polish rate near the edge of the wafer. This effect may be used

2.6.2 System Factors

The significant factors incorporated into Preston's coefficient, K_p, include:

1. Film type and properties
2. Abrasive particles, type, size, and morphology and concentration
3. Slurry compositon and pH
4. Temperature
5. Pad constitution, both bulk and surface structure (including conditioning effects).

2.6.3 Film Type and Properties

The dielectric films that are of primary interest in CMP are silicon dioxide films, grown or deposited by different processes. Other dielectric films that are polished include silicon nitride and silicon oxynitride. Polycrystalline silicon (poly-Si) is also considered with these films. Metal films are covered in Chap. 3.

In semiconductor technology, silicon dioxide films are used for many different applications. The CMP removal rate is a function of the specific process and operating point by which the silicon dioxide film is formed. Among these different technological approaches are low pressure chemical vapor deposition (LPCVD) and plasma enhanced chemical vapor deposition (PECVD). For each of these approaches, the reactants and operating points (temperature, pressure, ionizing energy, etc.) can vary widely. For one specific use, a region of operation of PECVD called high density plasma (HDP) is employed. It has become the preferred deposition process for shallow trench isolation (STI) structures. The structure and properties of silicon dioxide vary with process and operating point of the deposition process, and these, in turn, influence the CMP removal rate [25, 26]. Specifically, the film density and number of open bonds appear to correlate with CMP removal rate. Thermally grown silicon dioxide is the densest type film used in semiconductor processing, denser than deposited films, and polishes more slowly. This is pictured in Fig. 2.12. In Fig. 2.12, the doped (BPSG) films polish more quickly than do the undoped (USG) films. The denser HDP films polish more slowly than do the APCVD films, with thermally grown silicon dioxide, the densest film, polishing the most slowly. Silicon dioxide films doped with phosphorus and sometimes boron are widely used for the first dielectric layer covering the active devices. In current semiconductor production, these films are now planarized with CMP.

These first dielectric layer films (this level is sometimes referred to as ILD0) can contain varying amounts of boron and phosphorus. The CMP removal rate is a strong function of both dopants. Two graphs of CMP results picturing this dependence are shown in Fig. 2.13. Over the range of dopants tested, the removal rates increase with concentration of both dopants [27, 28].

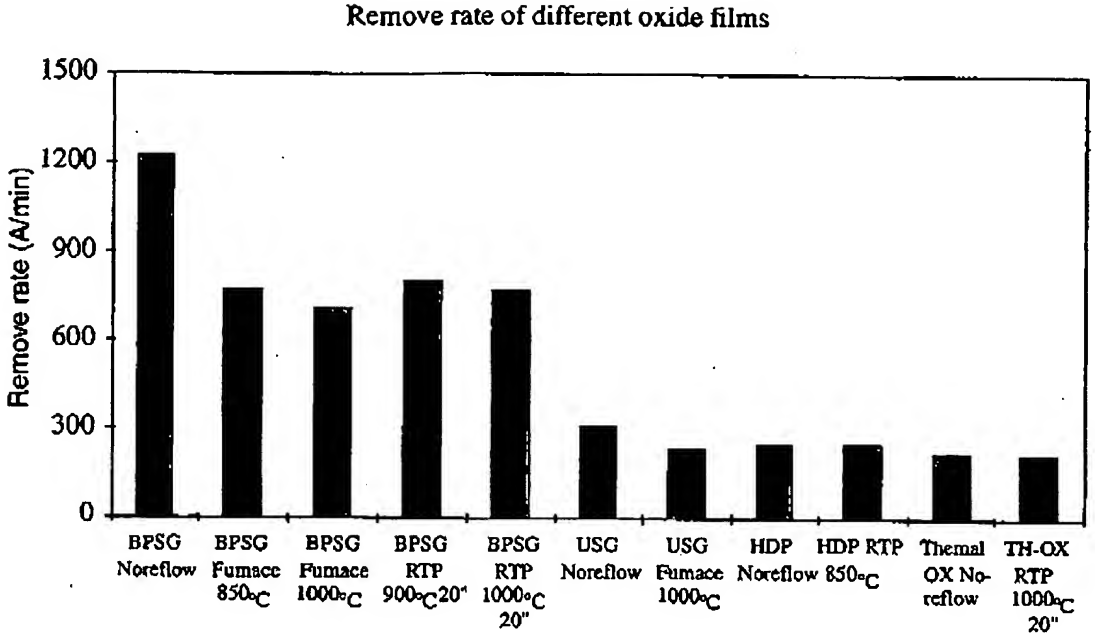

Fig. 2.12. Polish removal rates for different silicon dioxide films. Different deposition processes, dopant concentrations, and post-deposition anneals are compared, using thermally grown silicon dioxide as a reference. From [23]

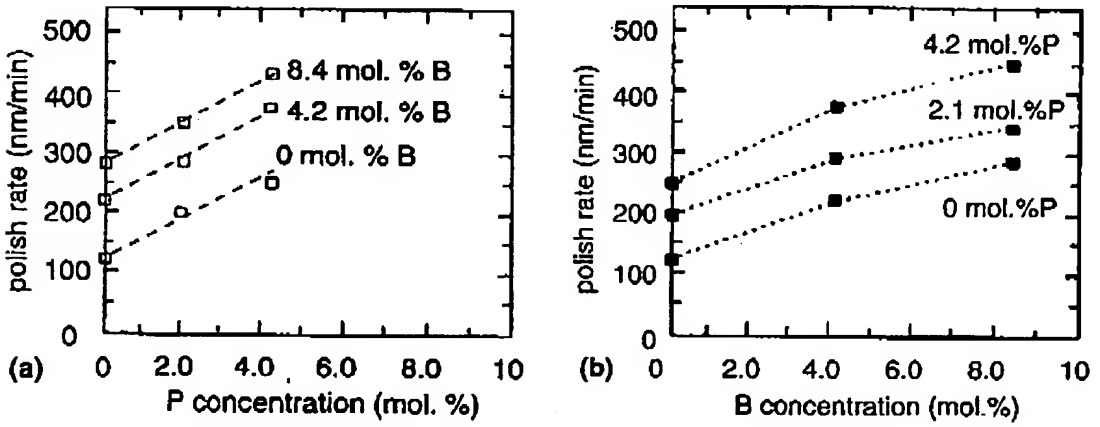

Fig. 2.13. (a) CMP rate of SiO_2 and BSG as a function of P concentration. (b) CMP rate of SiO_2 and PSG as a function of B concentration. From [25].

2.6.4 Abrasive Particles

The type, size, morphology and concentration of the abrasive particles in the slurry strongly influence the polish rate. If we consider initially only silica abrasive particles, there is a wide range of behaviors that are observed with changes in type, size, morphology and concentration. The primary types of silica abrasive used in CMP are fumed silica particles and colloidal silica particles. The fuming process [14] creates tightly bound aggregates of smaller

ticles are formed in solution and are, in general, nearly spherical, but the maximum particle size is usually smaller than can be achieved with the fuming process (see Chap. 7).

A typical silica abrasive slurry used in the industry is SS-12™ supplied by Cabot. The concentration of abrasive as well as other properties is listed in Table 2.1. As noted, the pH is near 11. The abrasive particles are created by a fuming process, which is described in Chap. 7 and in [14]. Fumed silica particles employed in consist of an aggregate of tightly bound primary particles about 20 nm diameter with the mean aggregate size in the range of 100–300 nm.

Table 2.1. Properties of Cabot Semi-Sperse 12 (SS-12™) silica slurry. Courtesy of Cabot Corp

Property	Value
pH	10.9–11.2
Viscosity (cps)	$\leq$15
Specific Gravity	1.071–1.078
Mean Aggregate Particle Size(nm)	130–180
% Solids	12.4–12.6

It was noted that abrasive particles are an essential component of the CMP system. Several researchers have shown that, in low concentrations and for other parameters held constant, that polish rate is linearly proportional to particle concentration in the slurry. At sufficiently high concentrations, the polish rate is sublinear with increasing particle concentration. For one type of particle, the colloidal silica used in 30N50pHN™, supplied by Rodel, Inc., the polish rate vs. particle concentration is shown in Fig. 2.14. As seen in the figure, the polish rate is linear with particle concentration up to about 20 weight %.

Particle size can also play a role, though in the range of particle sizes used in silica slurries, it does not appear to be a strong effect. For very small particles, the rate goes down for a given silica concentration as particle size is reduced. This is reviewed as well in Chap. 7.

2.6.5 Pad Conditioning

Pad conditioning is necessary to maintain the asperity structures on the surface of the polishing pad. The asperities on the pad surface force the abrasive particles against the wafer. The pad asperities need to be continually regener-

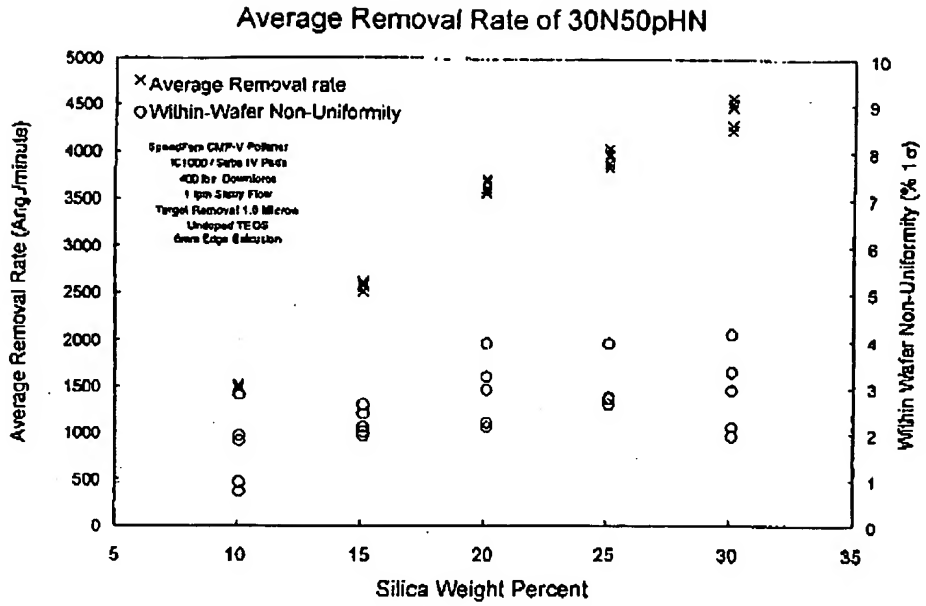

Fig. 2.14. Polish rate of silica slurry, 30N50pHN^{TM}, as a function of silica concentration in weight percent. Courtesy of Rodel, Inc.

done between wafer polishing cycles is called *ex-situ* conditioning. A representative graph of the reduction in polish removal rate when no conditioning is used to maintain the asperity profile is shown in Fig. 2.15, from [30]. The pad was conditioned normally between wafers (*ex-situ*) until this test was started. Then for this set of wafers no conditioning was done at all. Note that the polish rate decay is gradual and begins immediately when conditioning is not used.

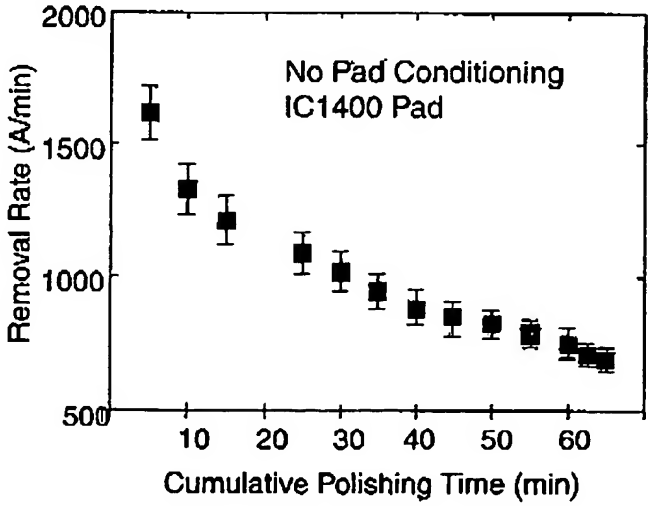

Fig. 2.15. Decrease of polishing rate in the absence of pad conditioning. From [27]

Because the removal rate decay begins immediately when conditioning stops, the latter part of any polishing step using an *ex-situ* conditioning process has a drop in the polishing rate during the polishing step itself [31]. The effective polish rate for a given step then is the average rate during the step and not the maximum rate. By using *in-situ* conditioning this problem is reduced, as the asperities are being generated simultaneously as they are being worn down. Some polishing cycles are 3 minutes long or more, and for these steps *in-situ* conditioning offers substantial throughput advantages.

The surface of a polishing pad is shown in Fig. 2.4. If one examines the land areas between the pore openings, the asperities created by conditioning this surface can be measured. This has been done for *ex-situ* conditioning where the polish rate over shourt intervals has been compared to the average asperity height measured on small coupons removed from the polishing pad [28].

The results for a standard IC1000^{TM} pad taken over eight one minute intervals, with no intermediate conditioning, are shown in Fig. 2.16a and 2.16b. The average removal rate for each interval as well as the average asperity height is shown for areas between the pore openings, where the asperities created by conditioning this surface can be measured. This has been done for *ex-situ* conditioning where the polish rate over short intervals has been compared to the average asperity height measured on small coupons removed from the polishing pad [28].

The results for a standard IC1000^{TM} pad taken over eight one minute intervals, with no intermediate conditioning, are shown in Fig. 2.16a and 2.16b. The average removal rate for each interval as well as the average asperity height is shown. As the asperity rate decreases, so does the polish rate. For representative polish conditions and rates, the asperity heights are in the 1–2 μm range. Note that the average asperity heights are much smaller than the average pore size of 30–50 μm.

2.6.6 Temperature and pH

Temperature and pH also affect the polish rate. As the pH increases in the regime near pH = 11, the polish rate increases. In practice, it is difficult to operate much above pH = 11.5, as the silica particles in the slurry begin to dissolve with time so that that the abrasive particles are not stable over time [15]. A representative curve of polish rate vs. pH is shown in Fig. 2.17. Also plotted is the polish rate vs. pH for CVD silicon nitride. From pH = 9.7 to pH = 10.7, the polish rate increases by about 20%. It is difficult to maintain silica in solution near and above pH = 11.5, so most silica slurries are made with pH near 11.

There is also an increase of silicon dioxide polish rate with ambient temperature near and above room temperature [26]. In [26], the authors attribute

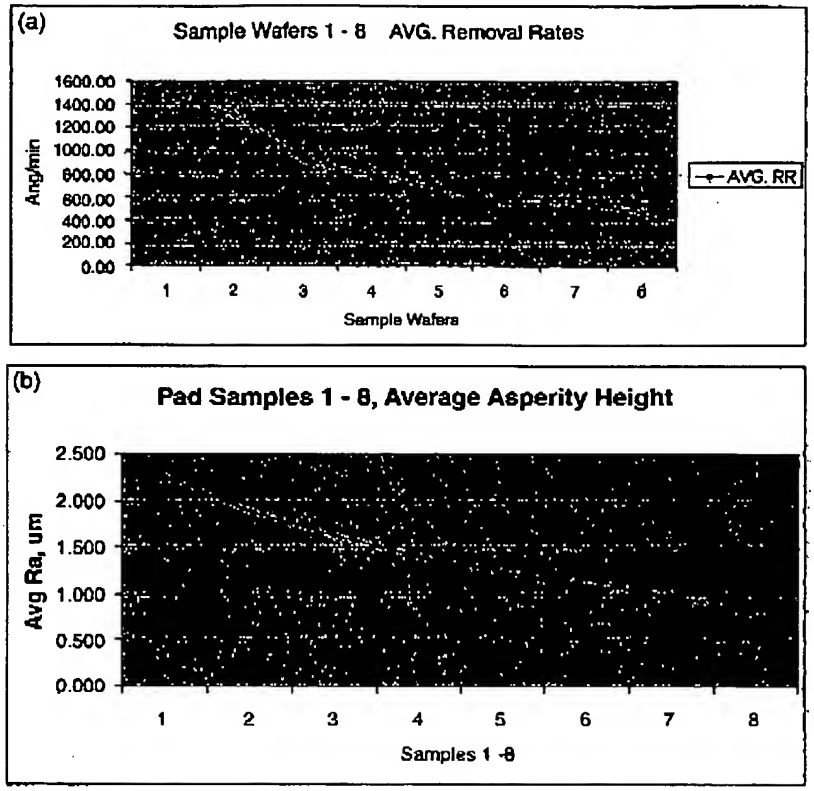

Fig. 2.16. (a) The average polish rate for 30 second intervals with no conditioning. (b) The average asperity height on the areas between the IC1000™ pores at the end of the 30 second intervals shown in Fig. 2.16a. The measurement was made by a Zygo NewView 5000. From [26]

higher temperatures. In practice, most polishing machines have systems to heat or cool the platens in order to optimize a given polishing process.

2.7 Scales and Random Polishing Effects

The major variables of the silicon dioxide CMP process have been discussed on an elemental scale. These variables are the factors that locally affect polish rate, including properties of the silicon dioxide film being polished as well as the pad and slurry properties.

On the smallest scale, as pictured in Fig. 2.5, the silicon dioxide film removal occurs when an abrasive particle is forced against the silicon dioxide film by an asperity, and the relative velocity of the asperity (on the pad) and the film creates a path of film removal. Global film removal occurs as this action is repeated a very large number of times. The directions of the polishing paths are randomized by having the film (on the wafer) rotate

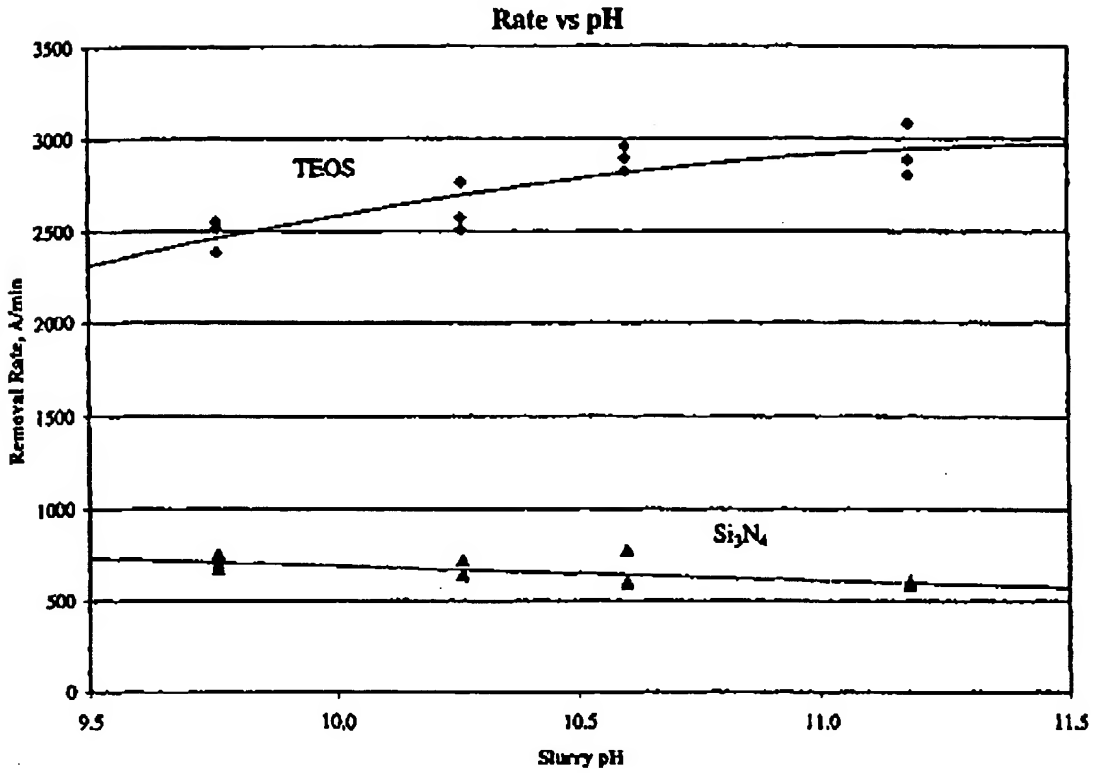

Fig. 2.17. Variation of CMP removal rate vs. pH for two films. Note that the TEOS (silicon dioxide) removal rate increases with increasing pH in this pH range

Local planarization is achieved when the higher points of the film surface topography are removed more quickly than the lower points. Global planarization issues associated with film pattern density variations have been discussed above. In addition to pattern density effects, there are several other issues that affect global planarization, especially on the wafer scale.

In Fig. 2.19, which is similar to Fig. 2.5, but on a somewhat larger scale. The dimensions of the features are near to scale. Metal lines are about 1 μm high, and the range of asperity heights is 1–3 μm. Here, the silica particles are pictured as irregular, but of course the shape and size distribution is determined by the manufacturing process (see Chap. 7). The asperity heights in a given process depend upon several parameters, with pad type and conditioner and conditioning process being the most influential.

If we look at the polishing system on a yet larger scale, features of the pad structure other than the asperity profile appear. In Fig. 2.20, the pores of an IC1000^{TM} pad are pictured and grooves are also shown. Both of these features enhance slurry flow between the wafer and the polishing pad (see Chap. 6). However, pads without grooves are sometimes used.

These local, random variations of the asperities and pore structure appear on a small scale. The statistical averages of these properties, such as the

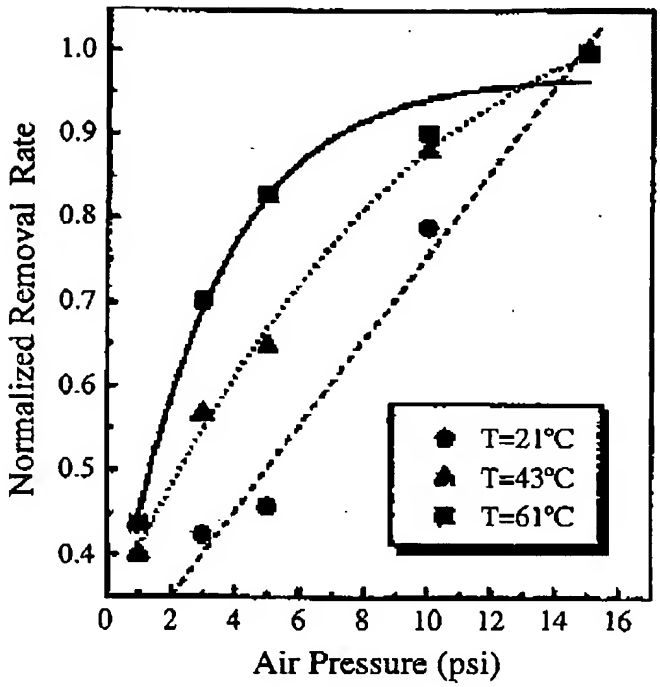

Fig. 2.18. The normalized oxide removal rate vs. air pressure used to provide the normal CMP load. Here, T is the slurry temperature, and the oxide removal rate for each slurry temperature is normalized to its highes value, respectively. In this experiment, an IC1000 pad was used. From [23]

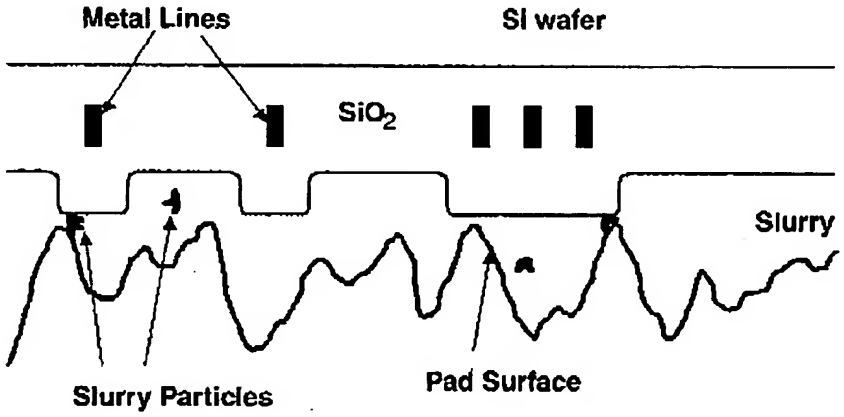

Fig. 2.19. Diagram of the elements of the CMP process showing a larger region than that of Fig. 2.5

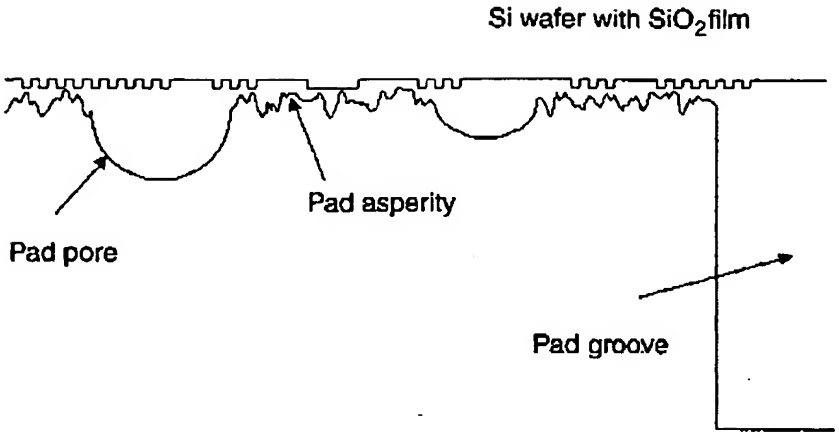

Fig. 2.20. Wafer with $\approx 2\,\mu m$ features with asperities of comparable dimensions and pores of 30–50 μm diameters. Also the edge of a pad groove is shown

Chap. 6), do affect the observed polishing behavior. This is the result of the averaging process of a great many polishing events over large areas of the polishing pad in all directions.

2.7.1 Polish Rate and Other Variations Introduced by the System

CMP polishing systems have matured to produce consistent, well controlled processes at all scales, within-die, within-wafer, and wafer to wafer. The general approach has been to provide as uniform polishing conditions as possible at all levels. Because of wafer and table geometries, both the down force and velocity applied to the wafer are not uniform across the wafer nor over time. To minimize the effect on the resultant polished film of these variations, table and carrier rotation are employed. In addition compressible elements, such as carrier films and bottom pads, are used to provide a more uniform applied down force. Looking at the wafer and system as a whole, forces are provided between the back of the carrier film and the platen surface. However, the polishing process takes place at **the interface between the pad surface and the film being polished.**

In polishing machines, the local velocity over time at the pad surface-film interface is very well determined by the geometry of the machine. The local pressure, in contrast, is sensitive to any lateral dimensional variations in the layers of materials between the ideal carrier head surface and the ideal platen surface. Lateral thickness variations over the wafer and over all the wafer paths traversed across the rotating platen will create a time and pattern dependent variation of pressure at the pad surface-film interface at any specific point or thereafter. In addition, if the elastic constants of the compressible elements, the carrier film and bottom pad, change with position or slowly with time as many wafers are polished, these compression changes

The methods to improve the uniformity of applied pressure at the pad surface-film interface have focused on machine design improvements (see Chap. 5), and on replacing pads and carrier films when within-wafer non-uniformity becomes too large during production CMP. Two major design approaches that have been implemented on new machines and sometimes retrofitted on older machines are 1) the position dependent conditioning which is designed to keep the thickness of the polishing pad uniform and not allow a trough to form in the track of wafer travel, and 2) the fluid backed wafer carrier, where the local pressure applied across the back of the wafer is by a fluid, applied either directly to the back of the wafer or through a thin membrane. Both of these improvements are discussed in Chap. 5.

A second issue is the edge effect, which is a strong variation in the polish removal rate as a function of radial position near the edge of the wafer. This pressure variation and the observed polish rate as a function of radial position are pictured in Fig. 2.21a and 2.21b from [28]. This effect can be reduced by varying the bottom pad stiffness and thickness as discussed by Baker. Note that, while carrier rotation can average out, to a large degree, the leading edge to trailing edge variation shown in Fig. 2.3, it will not affect the edge effect since the magnitude of the effect only has radial dependence.

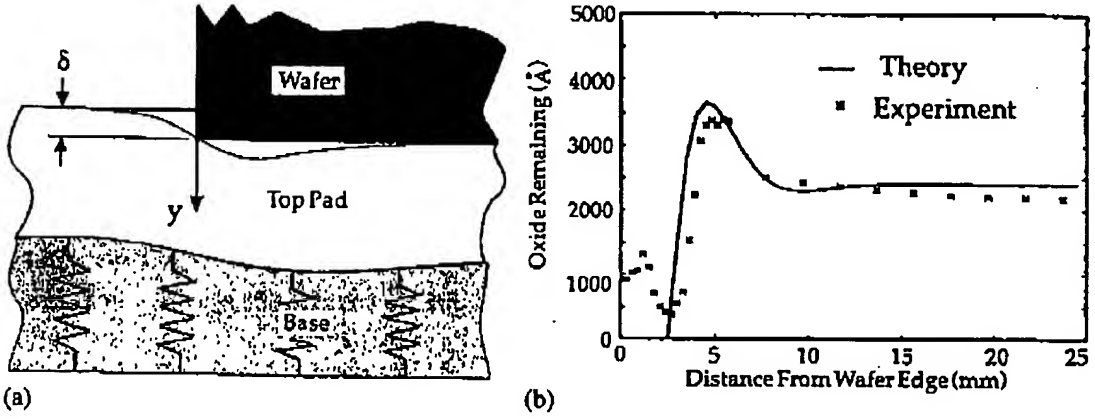

Fig. 2.21. (a) Model of pad structure that produces the edge effect. From [28]. (b) Model and data for remaining silicon dioxide at the edge of the wafer. From [28]

2.8 Random Effects

When we consider the entire wafer in the CMP system, there are several longer range mechanisms that can alter the polish rate across the wafer. One, discussed above, is the gimbal induced force that causes the leading

edge, resulting in higher polish rate on the leading edge as shown in Fig. 2.3. This higher pressure on the pad can be reduced by lowering the effective gimbal point of the carrier, or by avoiding the gimbal effect altogether by providing pressure uniformly [32] to the back of the wafer by a fluid under pressure.

In the initial description of the pads, wafer and carrier films, the soft bottom pad and carrier film were required to make the applied pressure at the wafer-spolishing pad interface more uniform. The soft pads and carrier film improve the pressure uniformity, but since they act as springs they do not eliminate it. Referring to Fig. 2.2, dimensional variations in all the elements (wafer, pads, and carrier film) and elasticity variations in the bottom pad and carrier film all can lead to local pressure variations across the wafer at the wafer-polishing pad interface.

The effects of the variations in elasticity of the bottom pad and the thickness variations of both pads are reduced because of the averaging effects of table and carrier rotations. However, carrier film thickness and elasticity variations and wafer thickness variations are not reduced by this averaging and so lead to local polished film variations.

2.9 Slurries with Particles Other than Silica

Though silica-based slurries are primarily used for CMP of silicon dioxide and other dielectrics, other particles can also be used. As noted earlier, ceria, CeO_2, is known from glass polishing experience to be much more effective than silica , in terms of polishing rate, at polishing a glass film. This naturally has led to its evaluation as a CMP polishing abrasive.

The data for polishing with ceria abrasives [33] shows that indeed that, per abrasive particle, ceria polishes planar surfaces much more effectively than does silica. As shown in Fig. 2.22, the polish rate for planar silicon dioxide surfaces is higher with a slurry containing 0.5 wt% ceria than for silica based slurries containing 13 wt% silica. Ceria behaves differently than does silica for structured surfaces however. For a structured surface such as an ILD surface, silica based slurries polish the high areas at a rate that is higher than the planar rate. For ceria based slurries, this effect is much less pronounced, and sometimes the initial polish rate can be lower than the rate for planar surfaces [33]. However, this effect appears to be sensitive to the type of ceria used, as well as to certain additives for some slurries. Ceria slurries are also very sensitive to the nature of the silica film being polished. Diffferent depostion conditions can lead to a very different polishing characteristics [35, 36].

Another abrasive that has been tested for polishing silicon dioxide is manganese sesquioxide, Mn_2O_3 [34]. However it has not been extensively investi-

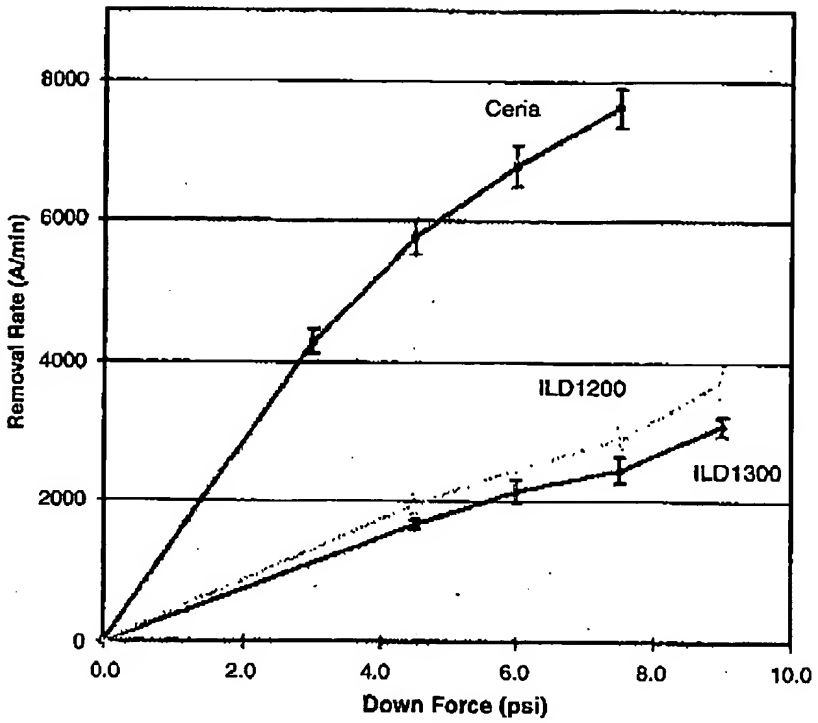

Fig. 2.22. Polish rates for two silica slurries with 13 wt% silica compared to a ceria slurry with 0.5 wt% ceria. The rotation rate is 40 rpm. From [33]

abrasive concentration is very non-linear, as shown in Fig. 2.23. This abrasive is used as a slurry in an alkaline pH range, but if the wafer is cleaned in an acidic solution, the residual particles are dissolved, thus simplifying the post-CMP cleaning of the wafers.

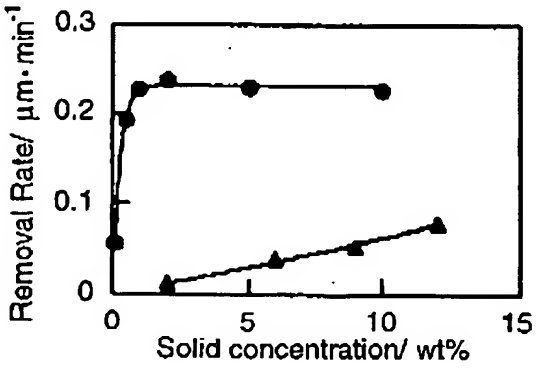

Fig. 2.23. Polish rate dependence of Mn_2O_3 slurry (●) compared to a silica slurry (▲). Note the very large polish rate as a function of concentration in the low concentration region. From [31] (©2003 IEEE)

2.10 Non-ILD Non-Metal CMP

2.10.1 Shallow Trench Isolation (STI)

The maturity of CMP technology aided the rapid introduction to semiconductor manufacturing of Shallow Trench Isolation (STI) technology [37]. As discussed in Chap. 10, this technology replaces LOCOS (LOCal Oxidation of Silicon) as an isolation approach. STI allows much tighter packing of transistors, thus increasing the number of transistors per unit area, with design rules being equal. As lithography dimensions, and their associated design rules, become smaller the relative advantage of STI vs. LOCOS becomes greater.

The standard approach to the CMP of an STI structure is shown in Fig. 2.24. The silicon substrate has a thin silicon dioxide buffer layer and a CVD silicon nitride layer on top of it. This stack is patterned and etched, with shallow trenches etched into the silicon substrate. The silicon nitride areas are where the active transistors will be placed. The structure is then filled with silicon dioxide, usually deposited with a technique called high density plasma (HDP). This deposition approach is able to fill very small trenches without leaving voids.

The CMP step is to remove all the silicon dioxide from the top of the silicon nitride, while polishing as little of the silicon nitride as possible. Then, in the overall process sequence the nitride is removed and the buffer silicon dioxide layer is etched off with hydrofluoric acid. The transistor gate oxide is then formed and polysilicon is then deposited as the gate electrode material.

CMP of the silicon dioxide-silicon nitride system has its unique considerations. When a standard silica based slurry is used, the silicon dioxide polish rate, for the same CMP conditions, is about three times greater than that of silicon nitride. This helps control the variation of the post-CMP thickness both within the die and within the wafer. Because of severe integration restrictions, as discussed in Chap. 10, a very tight control on the remaining silicon nitride thickness is critical. This is because the structure and shape of the edge of the transistor at the substrate-trench interface determine the electrical performance of the transistors. If there is much variability in the shape of the gate electrode edge, then the transistor performance and process yield will degrade.

The CMP step is to remove all the silicon dioxide from the top of the silicon nitride, while polishing as little of the silicon nitride as possible. Then, in the overall process sequence the nitride is removed and the buffer silicon dioxide layer is etched off with hydrofluoric acid. The transistor gate oxide is then formed and polysilicon is then deposited as the gate electrode material.

The behavior of CMP when polishing patterned wafers with variable density affects the STI process as it does the ILD process. Isolated raised features are polished quickly and dense areas are polished more slowly. This means

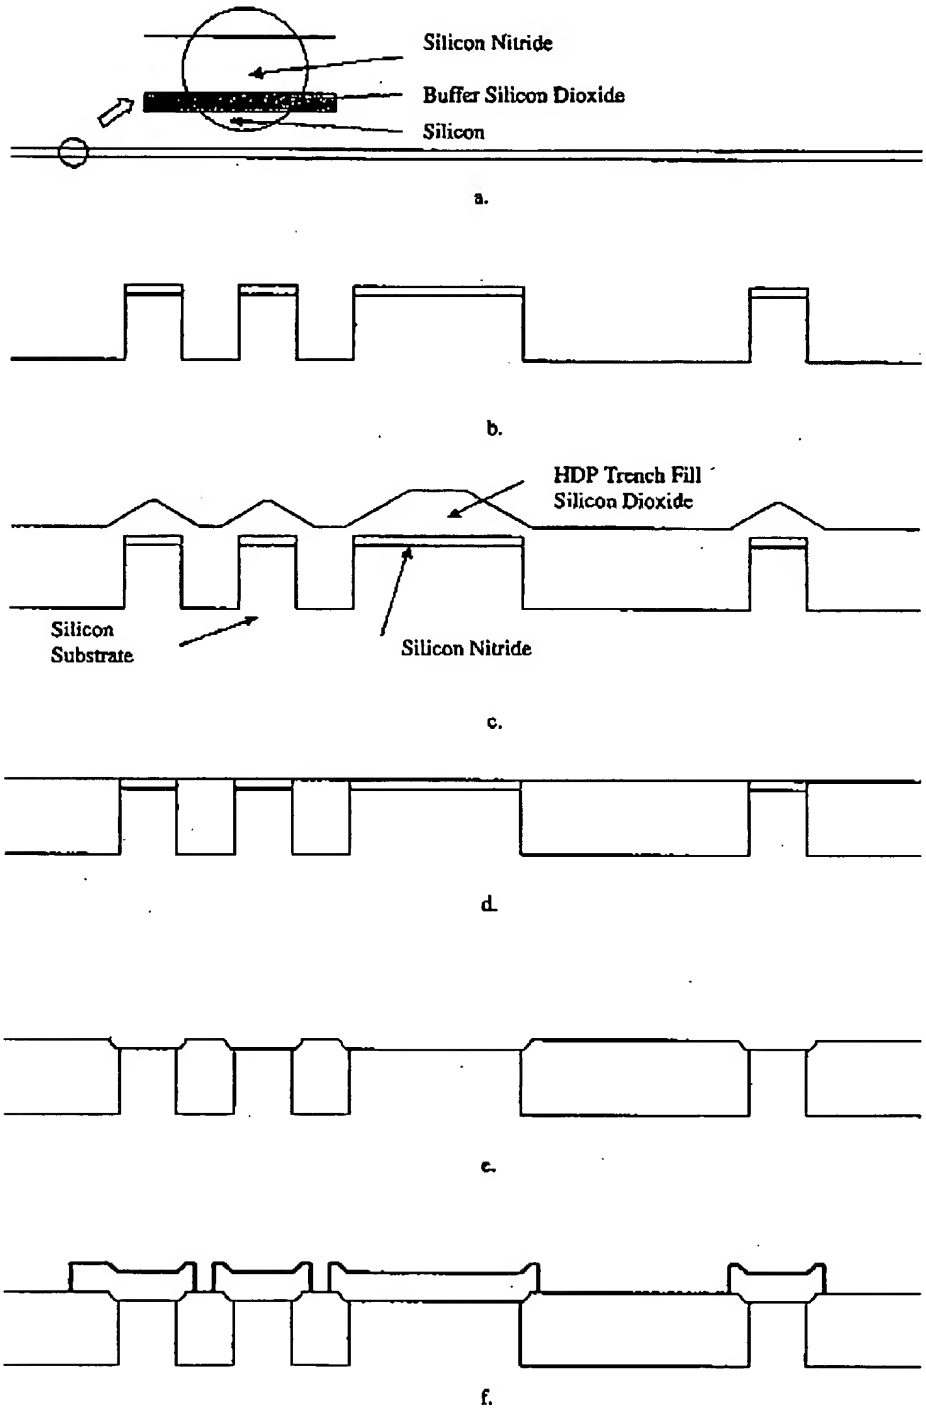

Fig. 2.24. Shallow Trench Isolation formation, including deposition and definition of poly Si gate. Starting with silicon nitride-buffer silicon dioxide films, the active areas are defined and the trenches are etched (a,b). After filling the trenches with silicon dioxide (here, HDP oxide), the oxide is planarized with CMP (c,d). Then the silicon nitride is removed, and the buffer silicon dioxide is also removed. Finally, the gate oxide is deposited and the poly-Si gates are formed (e,f)

that the nitride layer on an isolated feature is exposed and polished for longer times than is the nitride layer on a feature in a dense area.

Clearly, WIWNU and WIDNU of the remaining nitride and oxide thicknesses need to be controlled very tightly. The WIWNU issues are similar to those for ILD processing but the WIDNU issues are different. One major direction to minimize the WIDNU is to minimize the density variations so that, within a planarization length or so, all areas of the die have the same density pattern of silicon dioxide above the silicon nitride layer. In this way, the die itself will have less contribution to the variation in the time to clear the silicon dioxide from the silicon nitride.

Several approaches have been proposed to make the silicon dioxide density nearly uniform across the die. A reverse mask step is an approach that several groups have pursued [38]. This sequence is pictured in Fig. 2.25. Here, after the silicon dioxide is deposited, a reverse mask of the STI pattern, with features made slightly smaller than the STI mask itself, is patterned and the exposed silicon dioxide etched away. This leaves a fence of silicon dioxide around the edge of each feature but the total amount of oxide above the silicon nitride is sharply reduced.

A second approach is to make the density of STI structures as uniform as possible. This can be done by inserting dummy structures into the STI mask pattern. These structures are not to be anything more than isolated islands of conducting silicon. The dummy structure size and shape should approximate the active areas of most of the chip; usually this means minimize size transistors. This averaging by the use of dummy structures can only be approximate but it certainly can minimize the extremes of local pattern density variation [39, 40].

Another way to improve the performance of the CMP module is to make the slurry very selective. Slurry selectivity is the ratio of the polish rate of one film compared to another at the same CMP operating conditions. Here,

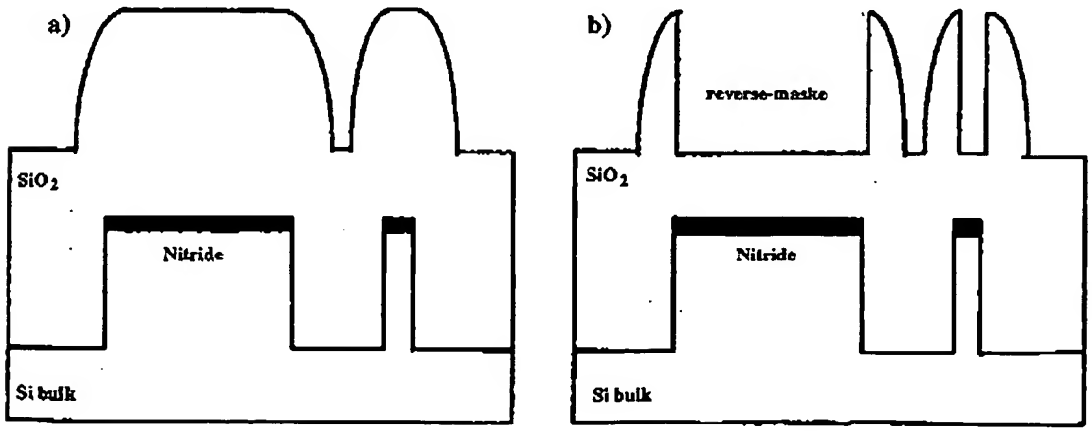

Fig. 2.25. The effect of reverse, or counter, mask etch back to reduce the amount

selectivity means rate selectivity, the ratio of the removal rates of silicon dioxide and silicon nitride. As noted above, for standard alkali-based silica slurries this selectivity is about 3:1. By using appropriate additives, very high selectivities, greater than 100:1, can be obtained [41]. These very high selectivity slurries can substantially reduce the non-uniformity of remaining silicon nitride across the die and wafer.

The use of high selectivity slurries is standard in metal CMP (see Chap. 3), and there are a number of unique effects that occur. Recess, or dishing, is one effect that occurs when the more slowly polishing layer is exposed. At that point the faster polishing layer continues to polish creating a dished area next to the slower polishing layer. For STI polishing this effect is very important and needs to be minimized [42].

Recess will affect the geometry of the gate electrode at the silicon-trench interface and needs to be controlled as tightly as the amount of nitride removal. Recess is influenced by the polishing operating conditions, most importantly the overpolish time required to clear the silicon dioxide from all the nitride structures [43, 44].

2.10.2 Polysilicon Polish

Silicon, including polysilicon, can be polished easily with basically the same types of polishers, and similar pads and slurries, that are used to polish silicon dioxide. In general, for the same polishing conditions, polysilicon (or polycrystalline silicon) deposited by LPCVD systems polishes more quickly than does silicon dioxide. However, with appropriate additives to the slurry, the rate selectivity between polysilicon and silicon dioxide can be made greater than 100:1 [45]. Using such slurries, several polysilicon CMP steps have been used. One is to polish polysilicon plugs, or vias, removing the layer of polysilicon on top of the ILD and leaving the plug filled with polysilicon [7, 46]. Another step that has been used is for polishing polysilicon to smooth it for subsequent processing [47].

In addition to these more widely used steps, there also have been integration sequences that have used polysilicon CMP in the formation of STI structures as well as interconnect structures.

2.10.3 Low K Dielectrics

With successive generations of semiconductor processed, the dimensions shrink but the materials also change. Copper is replacing aluminum because it has a lower resistivity. Also, dielectrics with lower permittivity, or dielectric constant (K), will replace silicon dioxide. Copper CMP and its associated issues are addressed in Chap. 3. The integration issues associated with different low K dielectrics are discussed in Chap. 10. In addition to the different low K

approaches for the low K dielectrics. Some of these approaches employ thin barrier layers which are used as etch stops and CMP stopping layers.

One of the first generation lower K dielectrics to be integrated into a production semiconductor process was fluorinated silica glass, or FSG [48, 21]). The permittivity of undoped silica is about 3.9, and by doping with fluorine a reduced permittivity of 3.5–3.6 can be achieved. There are reliability issues which limit the fluorine concentration in the silica. The CMP removal rates for these glasses are at or slightly above those of undoped silica for the same CMP process conditions.

The dielectric materials that are used for permittivities lower than that of FSG behave, in general, very differently from silica. There are new materials being developed continually to provide improved performance of those already available. Types of materials include spin-on homogeneous organic based materials, such as SiLK^{TM} [49] and silsesquioxanes [50, 51]. A second group of materials being investigated closely include the CVD deposited carbon-doped silicas, with Black Diamond^{TM} [52] and Coral^{TM} being two commercially available materials. Another large group of materials are the porous materials where pores, or voids, are contained in the bulk of the material [53, 54]. Materials with voids have produced very low permittivities, some below 2.0, but are generally mechanically weak and are difficult to polish.

Most of the integration approaches for incorporating these low permittivity materials into the back end structure use, as noted, stop layers and the CMP process does not see the low K material. However, all materials have to be robust enough not to degrade under the pressure of the polishing process, and also must maintain good adhesion to their surrounding materials. Adhesion of many of these low permittivity materials is a significant issue.

At present, copper is being integrated into semiconductor processes, and the integration of low permittivity materials is lagging in its implementation. This lag is largely due to the many difficulties that have been encountered in integrating any low permittivity material beyond FSG into a multi-level metal process.

2.11 Conclusion

CMP of silicon dioxide was the initial application and is still the largest application of CMP in the semiconductor industry. A majority of the characterization of CMP processes has been on silicon dioxide processes. For this reason, the initial focus of the book has used these processes as the baseline. In addition, the characterization and evaluation of the consumables, pads and slurries, as well as polishing tools, has focused on silicon dioxide polishing.

Almost from the beginning of its application to semiconductor processing, CMP has been used for metal polishing as well as for dielectrics. It has also become an enabling technology for the introduction of copper as an inter-

semiconductor technology. The treatment of both the technology and modeling of metal polishing build upon, but are quite different from their silicon dioxide CMP counterparts.

For the above reasons, the book has been organized with the silicon dioxide CMP technology as the first technology chapter. The following chapters have addressed metal CMP technology and models, and then the hardware of the CMP process, pads, slurries and polishing tools. Finally, other key areas are addressed, including topography evolution and modeling, post-CMP cleaning, and overall process integration.

The application of CMP in semiconductor processing was introduced in the late 1980's, and is now a critical technological component in driving improved chip performance. Only in the past few years, though, has the scientific basis for the technology begun to receive much interest. At present, the number of publications focusing on a detailed understanding is small but is growing quickly. It is hoped that this book will provide a background that will enable the reader to be able to read the current literature and understand the science and technology of CMP as it evolves in the near future.

References

1. M. Fury, "The Early Days of CMP", *Solid State Technology*, 81, May 1997.
2. D. Pramanik, V. Jain and K.Y. Chang, Proceedings 1991 VMIC Conference, 27, 1991.
3. D. Moy, M. Schadt, C.K. Hu, F. Kaufman, A.K. Ray, N. Mazzeo, E. Baran and D.J. Pearson, Proceedings 1989 VMIC Conference, 26, 1989; See also U.S. Patent 4,944,836, July 31, 1990.
4. G. Nanz and L.E. Camilletti, IEEE Trans. Semiconductor Mfg., 8, 382, 1995.
5. Cook, L. M., J. of Non-crystalline Solids, **120**, 152, 1990.
6. C.W. Kaanta, W.J. Cote, J.E. Cronin, K.L. Holland, P.I. Lee and T.M. Wright, IEDM Tech. Dig., 769, 1997.
7. C.W. Kaanta, W.J. Cote, J.E. Cronin, H.S. Landis, W. Hill and J. Ryan, Proceedings 1991 VMIC Conference, 144, 1991.
8. S.S. Cooperman, xxx, Journal Electrochem. Soc., **9**, 3180, 1994.
9. S.D. Hosali, A.R. Sethuraman, J-F. Wang, L.M. Cook and D.R. Evans, Proceedings 1997 CMP–MIC Conference, 52, IMIC, Tampa, 1997.
10. K. Itabashi, S. Tsuboi, H. Nakamura, K. Hashimoto, W. Futoh, K. Fukuda, I. Hanyu, S. Asai, T. Chijimatsu, E. Kawamura, T. Yao, H. Takagi, Y. Ohta, T. Karawawa, H. Iio, M. Onoda, F. Inoue, N. Nomura, Y. Satoh, M. Higashimoto, M. Matsumiya, T. Miyabo, T. Ikeda, T. Yamazaki, M. Miyajima, K. Watanabe, S. Kawamura and T. Taguchi, 1997 Symposium on VLSI Digest, 21, 1997.
11. F. Preston, J. Soc. Glass Technology, **11**, 214, 1927.
12. D.J. Stein and D.L. Hetherington, Electrochemical Soc. Proc., **99–37**, 217, 1999.

14. D.R. Evans and M.R. Oliver, Proc. MRS, San Francisco, 2001.
15. P. Truong and L. Blanchard, Proceedings 1998 CMP–MIC Conference, 351, IMIC, Tampa, 1998.
16. H. Muelenweg, F. Klaessig, W. Lortz, G. Varga and A. Gutsch, Proceedings 2000 CMP–MIC Conference, 325, 2000.
17. R.K. Iler, *The Chemistry of Silica*, 42, Wiley, New York, 1979.
18. Y.C. Shih, E.M. Shamble, J-F. Wang, A.R. Sethuraman, H-M. Wang, R.L. Lavoie and L.M. Cook, Proceedings 1997 CMP–MIC Conference, 237, IMIC, Tampa, 1997.
19. B. Zhao and F.G. Shi, Proceedings 1999 CMP–MIC Conference, 13, IMIC, Tampa, 1999.
20. J-Z. Zheng, V. Huang, M. Toh, C. Tay, F. Chen and B.B. Zhou, Proceedings 1997 CMP–MIC Conference, 315, IMIC, Tampa, 1997.
21. D.R. Evans, B.D. Ulrich and M.R. Oliver, Proceedings 1998 CMP–MIC Conference, 347, IMIC, Tampa, 1998.
22. M. Weling, S. Bothra, C. Drill and C. Gabriel, Proceedings 1997 CMP–MIC Conference, 65, IMIC, Tampa, 1997.
23. B. Stine, D. Ouma, R. Divecha, D. Boning, J. Chung, D.L. Hetherington, I. Ali, G. Shinn, J. Clark, O.S. Nakagawa and S-Y. Oh, Proceedings 1997 CMP–MIC Conference, 266, IMIC, Tampa, 1997.
24. J. Grillaert, M. Meuris, N. Heylen, K. DeVriendt, E. Vrancken and M. Heyns, Proceedings 1998 CMP–MIC Conference, 79, IMIC, Tampa, 1998.
25. C.N. Huang, H.B. Liu and J.T. Lin, Proceedings 1998 VMIC Conference, 632, IMIC, Tampa, 1998.
26. I.J. Malik, T. Mallon, B. Withers, R. Emami, D. Mordo and I. Goswami, Proceedings 1997 CMP–MIC Conference, 209, IMIC, Tampa, 1997.
27. M.A. Jaso, J.P. Gambino, K. Huckels, M. Ilg and G. Coleman, Proceedings 1997 CMP–MIC Conference, 15, IMIC, Tampa, 1997.
28. M.R. Oliver, R.E. Schmidt and M. Robinson, Electrochemical Soc. Proc., **2000–26**, 77, 2000.
29. C.W. Liu, B.T. Dai and C.F. Yeh, Thin Solid Films, **270**, 607, 1995.
30. D.L. Hetherington and K. Achuthan, Semicon West CMP Symposium Proceedings, San Francisco, July 1998.
31. A.R. Baker, Electrochemical Soc. Proc., **96–22**, 228, 1996.
32. T. Osterheld, S. Zuniga, S. Huey, P. McKeever, C. Garretson, B. Bonner, D. Bennett and R.R. Jin, Proc. MRS Symposium, **566**, 63, 2000.
33. S-I. Lee, C-I. Kim, H. Kim, J-H. Kim, C-W. Nam, S. Kim and C-T. Proceedings 1997 CMP–MIC Conference, 163, IMIC, Tampa, 1997.
34. S. Kishii, R. Suzuki, A. Ohishi and Y. Arimoto, IEDM Tech. Dig., 465, 1995.
35. M.R. Oliver, D.R. Evans, D.L. Hetherington, D.J. Sten, J.E. Stevens and S.D. Hosali, Proceedings 1999 CMP–MIC, 383, IMIC, Tampa, 1999.
36. D.J. Stein, D.L. Hetherington, M.R. Oliver, S.H. Hosali, D.R. Evans and B. Her, MRS, Paper M3.8, San Francisco, 2001.
37. J.T. Pan, D. Ouma, P. Li, D. Boning, F. Redeker, J. Chung and J. Whitby, Proceedings 1998 VMIC Conference, 467, IMIC, Tampa, 1998.
38. M. Jouty, M. Rivoire and T. Detzel, Proceedings 1999 CMP–MIC Conference, 329, IMIC, Tampa, 1999.

40. B. Lee, D. Boning, D.L. Hetherington and D.J. Stein, Proceedings 2000 CMP–MIC Conference, 255, IMIC, Tampa, 2000.
41. T. Detzel, S. Hosali, A. Sethuraman, J-F. Wang and L. Cook, Proceedings 1997 CMP–MIC Conference, 202, IMIC, Tampa, 1997.
42. V.S.K. Lim, F. Chen, W.L. Goh, A. See, C.H. Loh, C. Lin, Q.H. Zhong and M. Xin, Proceedings 2000 CMP–MIC Conference, 177, IMIC, Tampa, 2000.
43. B. Withers, E. Zhoa, R. Jairath and S. Hosali, Proceedings 1998 CMP–MIC Conference, 319, IMIC, Tampa, 1998.
44. F. Chen, C. Tay, B.B. Hou, F.L. Chin and J-Z. Zheng, Proceedings 1998 VMIC Conference, 491, IMIC, Tampa, 1998.
45. S. Fang, R.B. Knamankar, B.B. Shinn, F. Abbasi and F. Zhang, Proceedings 1998 CMP–MIC Conference, 134, IMIC, Tampa, 1998.
46. G.H. Koh, xxx, Proceedings 1998 CMP–MIC Conference, 15, IMIC, Tampa, 1998.
47. M. Ravkin, J. Zhang, K. Mikhaylich, D. Hetherington and D. Stein, Proceedings 1999 CMP–MIC Conference, 297, IMIC, Tampa, 1999.
48. C.P. Chen, C.T. Lee, C.F. Lin, H.C. Yung and L. Fang, Proceedings 1996 CMP–MIC Conference, 82, IMIC, Tampa, 1996.
49. C.L. Borst, W.N. Gill and R.J. Gutmann, Proceedings 1999 CMP–MIC Conference, 409, IMIC, Tampa, 1999.
50. K. Barla, O. Demolliens, C. Gounelle, C. Lair, Y. Lafarges, V. Lasserre, S. Lis, E. Louis, C. Maddalon, Y. Morand, G. Passamard, F. Pires and C. Verove, Proceedings 1998 VMIC Conference, 25, IMIC, Tampa, 1998.
51. H.D. Jeong, H.S. Park, H.J. Shin, B.J. Kim, H.K. Kang and M.Y. Lee, Proceedings 1999 International Interconnect Technology Conference, 190, 1999.
52. M. Naik, S. Parikh, P. Li, J. Educato, D. Cheung, I. Hashim, P. Hey, S. Jenq, T. Pand, F. Redeker, V. Rana, B. Tand and D. Yost, Proceedings 1999 International Interconnect Technology Conference, 181, 1999.
53. C. Jin and J. Wetze, Proceedings 2000 International Interconnect Technology Conference, 99, 2000.
54. E.T. Ryan, H-M. Ho, W-L. We, P.S. Ho, D.W. Gidley and J. Drage, Proceedings 1999 International Interconnect Technology Conference, 187, 1999.